

# DoD High Performance Computing Modernization Program: a Distributed Computing Program

January 2015

www.hpc.mil



**Larry Davis** 

## HPC MODERNIZATION PROGRAM

## **HPC Modernization Program**

#### **VISION**

A pervasive culture existing among DoD's scientists and engineers where they routinely use advanced computational environments to solve the most demanding problems transforming the way DoD does business—finding better solutions faster.

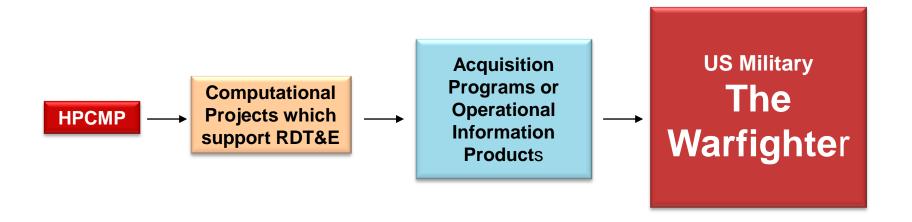
#### **MISSION**

Accelerate development and transition of advanced defense technologies into superior warfighting capabilities by exploiting and strengthening US leadership in supercomputing, communications, and computational modeling.



## **Technology Transition to the Warfighter**

#### **HPCMP Simplified Value Chain**



### **5 Execution Priorities**



#### Maintain Leadership

- Continue to effectively deliver high performance computing capabilities to the RDT&E community
- Pursue new service delivery models that better align the services with unique customer requirements
- Assume a leadership role in green high performance computing

#### Expand the Base

- Expand the application of HPC into emerging and non-traditional DoD mission areas
  - Acquisition/engineering communities, cyber, medicine, big data....

#### Increase Customer Productivity

- Provide portal tools that will support the seamless execution of workflows, make HPC services easily consumable by the next-generation workforce, and address the security barriers associated with connecting distributed desktops to centralized resources
- Keep pace with available cloud services



### 5 Execution Priorities (continued)

#### Exploit Next-Generation Technologies

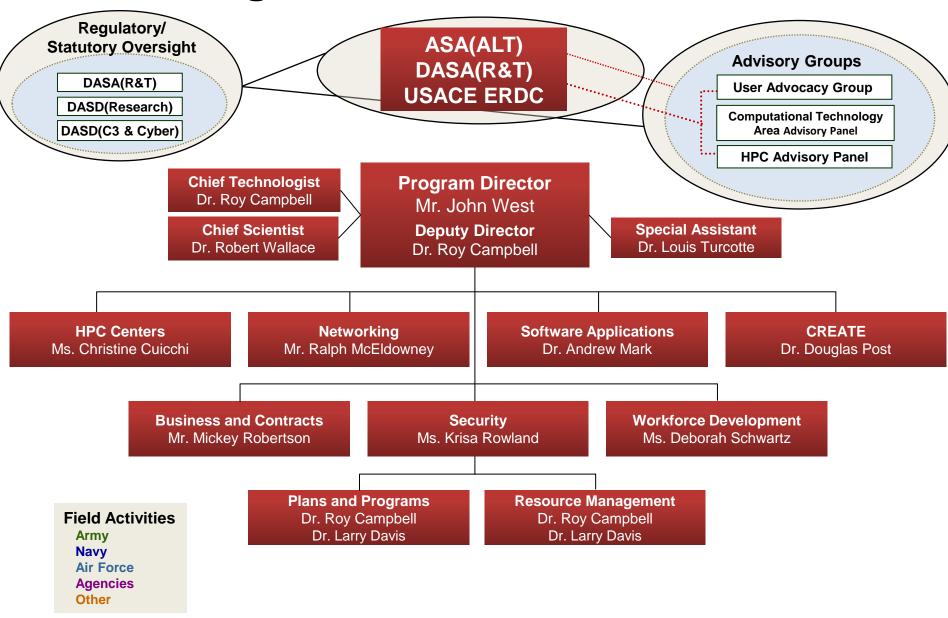
- Anticipate and exploit next-generation HPC technologies
  - Involvement in exascale initiatives, research investments in future HPC capabilities
  - Development of a software investment strategy that will position the DoD to efficiently consume future computing and communications capabilities

#### Develop the Workforce

- Develop and provide materials and experiences that create both a computationally literate technical workforce for the DoD and a cadre of computational professionals to provision HPC services
  - Investments in STEM programs, support for internships, and involvement in national initiatives that attract and engage with the next-generation workforce

## **HPCMP** Organization Chart







## Service/Agency Approval Authorities (S/AAAs)

- Allocate Service/Agency controlled CPU hours on HPCMP systems
- Provide coordination with customer organizations, the HPCMP, and other S/AAAs
- Ensure Project Leaders and authorized users are performing work in support of DoD
- Implement requirements surveys, resource allocation, resource monitoring and resource reallocation
- Provide guidance to users on which HPC assets are appropriate for their projects
- Trade allocations with other S/AAAs

## HPCMP Provides a Complete Set of Integrated Advanced Computing Capabilities

HPC MODERNIZATION PROGRAM

- Developing next-generation scientific and engineering software for advanced computing
  - Fostering similar development by other DoD organizations
- Advanced computing capabilities via five Centers and an integrated network
  - Efficiently: Costs are 2–5 times less expensive than commercial cloud solutions for similar capability
  - Effectively: HPCMP resource utilization averages approximately 90% of available capacity
- Fostering the development of a computational workforce
- Integrating these capabilities to solve complex DoD problems

## POPPER MODERNIZATION PROGRAM

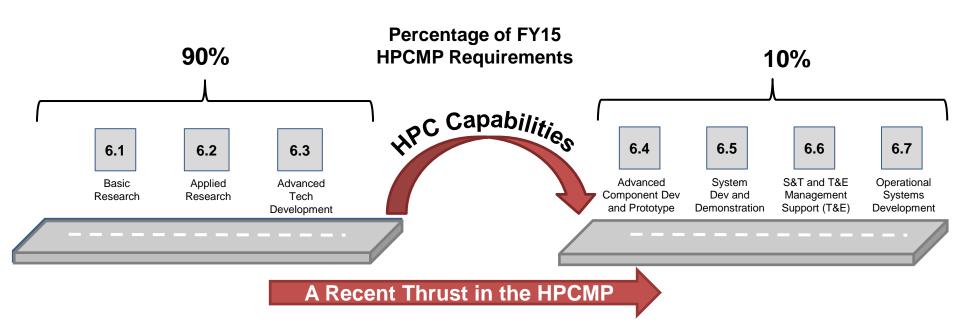
## What the HPCMP Offers Currently

- Approximately a dozen large HPC systems from a variety of vendors at five DoD Supercomputing Resource Centers
- Connectivity to these centers via a capable wide-area network the Defense Research and Engineering Network
- Expert help in all aspects of setting up accounts and using HPC systems through the Consolidated Customer Assistance Center and the User Productivity Enhancement, Technology Transfer and Training program
- Access to many commercial applications software packages for analysis
- Access to CREATE design modeling tools as they are developed

All members of the DoD RDT&E community are eligible to access and use these capabilities at no cost to the end user



## Percentage of Total FY15 Requirements by Budget Area





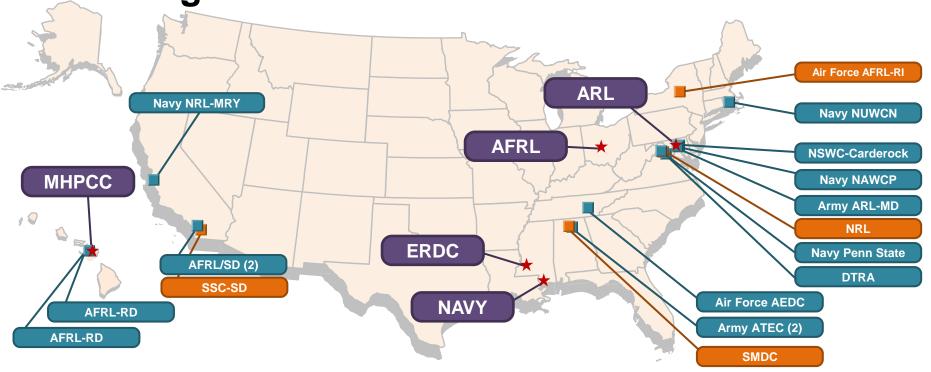


		DOD Projects		
Computational Technology Area	Leads	Primary	Not Primary	
Computational Structural Mechanics (CSM)	Carl Dyka, NSWCD	88	180	
Computational Fluid Dynamics (CFD)	Roger Strawn, AMRDEC	226	295	
Computational Chemistry, Biology, and Materials Science (CCM)	William Mattson, ARL	101	134	
Computational Electromagnetics and Acoustics (CEA)	Andrew Greenwood, AFRL	49	88	
Climate/Weather/Ocean Modeling and Simulation (CWO)	William Burnett, NAVO	39	51	
Signal/Image Processing (SIP)	Paul Antonik, AFRL	23	57	
Forces Modeling and Simulation (FMS)	Jennifer Park, SPAWAR	15	39	
Environmental Quality Modeling and Simulation (EQM)	David Richards, ERDC	2	11	
Electronics, Networking, and Systems (ENS)/C4I	VACANT	14	29	
Integrated Modeling and Test Environments (IMT)	VACANT	14	30	
Space and Astrophysical Sciences (SAS)	Jill Dahlburg, NRL	8	19	
Other		16	26	
Grand Total		595	959	

<sup>\*</sup> FY 2015 requirements survey



**HPCMP Supercomputing Resources Delivering HPC to the User** 



**DOD Supercomputing Resource** Centers (DSRC)

- US Air Force Research Lab (AFRL)
- Maui HPC Center (MHPCC)
- US Army Engineer Research and Development Center (ERDC)

**Affiliated Resource Centers (ARC)** 

- US Air Force Research Lab, Information Directorate, AFRL-RI Naval Research Lab (NRL)
- US Army, Space and Missile Defense Command SMDC
- US Navy, SSC-SD

NAVY

DTRA

• US Army Research Lab (ARL)

**Dedicated HPC Project Investments** (DHPI) FY11/12/13

- US Navy, Penn State
- US Army, ATEC (2)
- ¥US Air Force, AEDC
- US Air Force, AFRL/RD
- US Air Force, AFRL/SD (2)
- US Army, ARL-MD
- US Navy, NSWC-Carderock
- ŸUS Navy, NAWCP

¥US Navy, NRL-MRY

#### **Five Centers**

## MODERNIZATION PROGRAM

### Roles and Responsibilities

- Classified Computing Services
- Unclassified, but Sensitive, Computing Services
- CDAAC
- CDR

 Unclassified, but Sensitive, Computing Services

- HPC Helpdesk
- SW Configuration Management
- IAM Panel Lead
- ESM/ECM

 Unclassified, but Sensitive Computing Services

- Open Research Computing Services
- UDAAC
- UIT/ezHPC and pIE
- WMS

ARL



**AFRL** 



**ERDC** 



- Unclassified, but Sensitive, Computing Services
- Classified Computing Services
- UDR
- 24x7 Oceanographic Operations Modeling

**NAVY** 



- Unclassified, but Sensitive, Computing Services
- Special Computing Services
- · UIT Client Development
- Portal Development

**MHPCC** 



Strength Through Shared Services & Leadership



## DOD Supercomputing Resource Centers (DSRCs)

## **Current Capability**

Location	System	Available Processors (Cores and Accelerators)	Memory <i>(GB)</i>	Peak Performance ( <i>TeraFLOPS</i> )
AFRL	SGI ICE X	73,440	146,880	1,528
	SGIUV	1,004	3,514	22
	Cray XC30	56,880+(32 GPU)	153,088	1281
	SGI ICE X	125,888+(356 Phi)+(356 GPU)	453.136	5620
ARL	Cray XT5 (Classified)	15,600	41,600	100
	IBM iDataPlex	20,160	45,696	419
	IBM iDataPlex (Classified)	17,472	69,888	363
	Cray XC40	101,312+(32 GPU)	405,248	3774
ERDC	Cray XE6	150,912	301,824	1509
	Cray XE6 (Open Research)	14,976	29,952	138
	SGI ICE X	125,440+(32 GPU)	443,584	4662
MHPCC	IBM iDataPlex	12,096	24,192	253
NAVY	IBM iDataPlex (Classified)	4,032	8,960	84
	IBM iDataPlex	19,520	40,064	407
	IBM iDataPlex	19,520	40,064	407
	Cray XC30	29,352+(124 Phi)	82,752	786
	Cray XC30	28,584+(124 Phi)+(32 GPU)	82,112	822
	Cray XC40	50,208+(168 Phi)	200,000	1982
	Cray XC40	50,208+(168 Phi)	200,000	1982
	Cray XC40 (Classified)	6,720+(24 Phi)	27,109	266
	Total	923,324	2,346,980	26,405

As of: December 2014

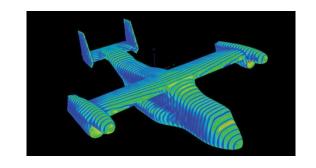
FY15 HPC Systems shown in PURPLE (TI-15 systems operational in Q4 FY15)
FY15 HPC Systems shown in GREEN (TI-14 systems operational in Q3 FY15)
FY14 HPC Systems shown in LIGHT BLUE
FY13 HPC Systems shown in ORANGE
FY11 HPC Systems shown in BLUE (TI-11/12 HPC systems operational in FY12)
FY09 HPC Systems shown in RED



## HPC Centers' Technical Services Support

#### Technical Services

- Customer Assistance and Support
- System and Network Administration
- System Installation and Integration
- Application/System Performance Analysis
- Data Analysis and Assessment
  - Classified and unclassified capability
- Operators and Operations Support
  - Enterprise System Monitoring (ESM)
- Outreach







## Consolidated Customer Assistance Center (CCAC)

- First point-of-contact for HPCMP user
- Telephone, user-portal, and email support
- Single ticket tracking system across all sites
- Tier 0 Self Help/Knowledge
   Management System on CCAC User
   Portal
- Tier 1 and some Tier 2 support CCAC
- Tier 2 support Centers
- Tier 3 support OEMs, PETTT, & Centers
- 11,883 tickets processed in FY14



Ticketing System

**CCAC** 

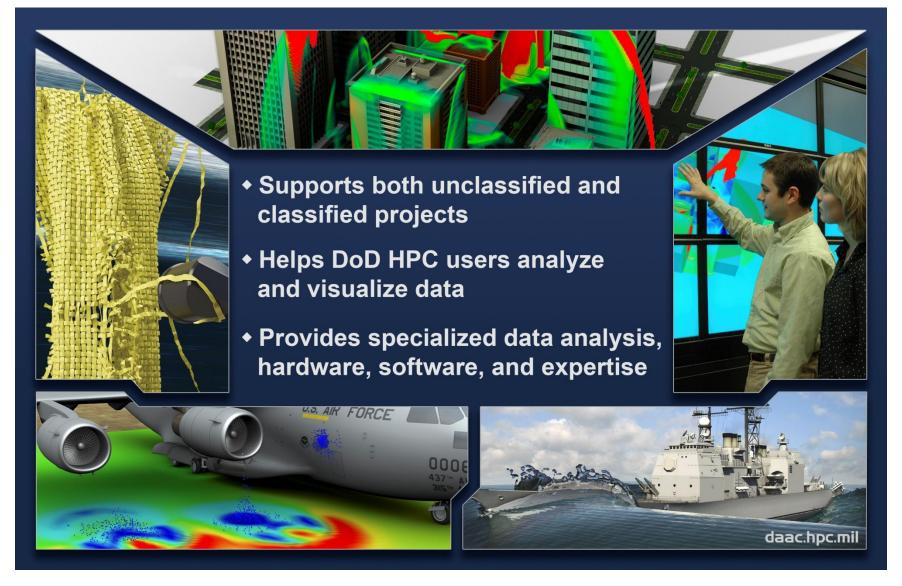
- Help Desk (Tier 1)
- Tier 2

CCAC Liaisons@ Centers

- Tiers 2
- Vendors 3



### **Data Analysis and Assessment Center**







## Operating Policies to Make System Diversity User Friendly

#### **Common Software Environment:**

Open-Source Libraries, Utilities &Compilers, Editors, Scripting Tools, Environment Variables, Debuggers, Login Shells, Multi-version SW Modules

#### **Standardized Processes:**

Queue Names, User Notices, Security Access

#### **User Support Features:**

Cross-Center File Transfer, Sample Code Repository, Ticket Life Consistency





## DOD Supercomputing Resource Centers Highlighted Projects/Services

#### User Dashboard

Web display of user-specific metrics and real-time system information

#### Software Configuration Management

Sharing software licenses across multiple Centers

#### HPC Portal

Desktop access to HPC

#### Advance Reservation Service

Reserve CPUs for a specific start time

#### Dedicated Support Partition

Dedicated CPUs for an extended period

#### Storage Management

- Data management for the customer
- Data Recovery

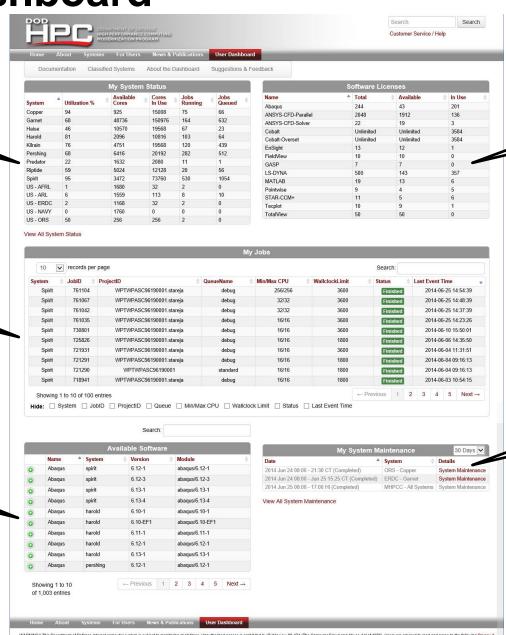




System Status

Job Status & History

Available Software



Security Notice, Date of Last Update: January 30, 2014

Software Licenses

System Maintenance

MB Revised: 7/1/2014



## **Software Configuration Management Sharing Licenses Across the Enterprise**



ARL DSRC



AFRL DSRC



NAVY DSRC



MHPCC DSRC



ERDC DSRC

Package	Baseline as of January 2015		
Abaqus	Analysis tokens: 244 & add-ons; 26 & add-ons (classified);		
Accelrys	Site licenses: MS_COMPASS, MS_DISCOVER, MS_DMOL3-SOLID-STATE, MS_CASTEP_PARALLEL; Flexible tokens: 8; MS_VISUALIZER 8;		
ANSYS	Negotiated WAN model: 23 CFD Solvers, 2048 CFD HPC, 5 Gambit (locked to geographic pools); 1 CFD Solver, 192 CFD HPC (classified)		
CFD++	2 Site licenses: ARL, AFRL; Floating: 2 Serial + 60 Parallel; Add on Modules: EDP & TNE		
Cobalt	Site licensed at all sites – Currently in procurement process		
EnSight Suite	13 Standard, 10 Gold, 5 EnLiten Gold, 2 DR 6 Standard, 2 DR (classified)		
GASP	7 (unlimited cores per job)		
Gaussian Suite	3 site licenses: 1 ARL, 1 AFRL, 1 ERDC		
LS-DYNA & 3D	500		
MATLAB	21 (3 Classified), 64 Dist Comp Toolbox, 10 Parallel Comp Toolbox		
Star-CD, Star-CCM+	11 Power Session licenses (unlimited cores per job)		
TotalView	50 architecture-free		
DDT	8,260 maximum processes		
FieldView	10 Utility Servers, 5 Standard and 5 batch AFRL		
Pointwise	5 Utility Servers, 4 AFRL. 4 ARL		
Tecplot	10 Utility Servers, 9 AFRL, 8 ARL		
Utility Server Software			
INTEL	5 Cluster Toolkit seats		



## **HPC Portal – Web Supercomputing**

#### **Quick Access to HPC**

#### Easy to Use

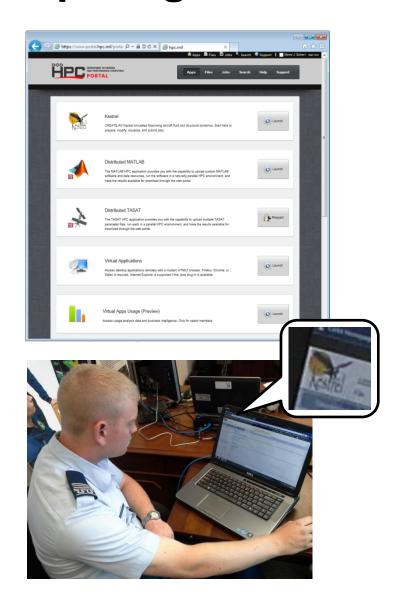
- Similar to a webmail interface
- No user-installed software or patches
- Integrated videos, forums, and help

#### Secure

- Quick DOD CAC-authentication
- Secured at one server vs. many desktops
- YubiKey enabled for University and Industry collaborators
- No desktop install is a security best-practice

#### Powerful

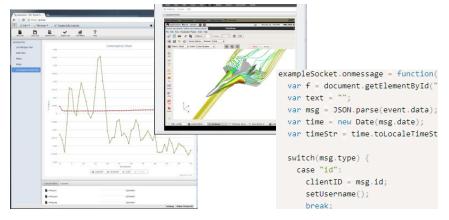
- Access to more than 50,000 CPU Cores
- Power-user command-line shell access
- Many apps at one location
- Software runs close to large datasets



#### **HPC Portal Benefits**



## **Applications for DOD-Scale Solutions**



#### **Productivity via Workflows**



Integrated analysis and engineering

#### HPC applications

- CREATE software
- Domain-specific tools
- General analysis and visualization
- Software Developer Kit to build ecosystem

#### **Collaboration and Learning**

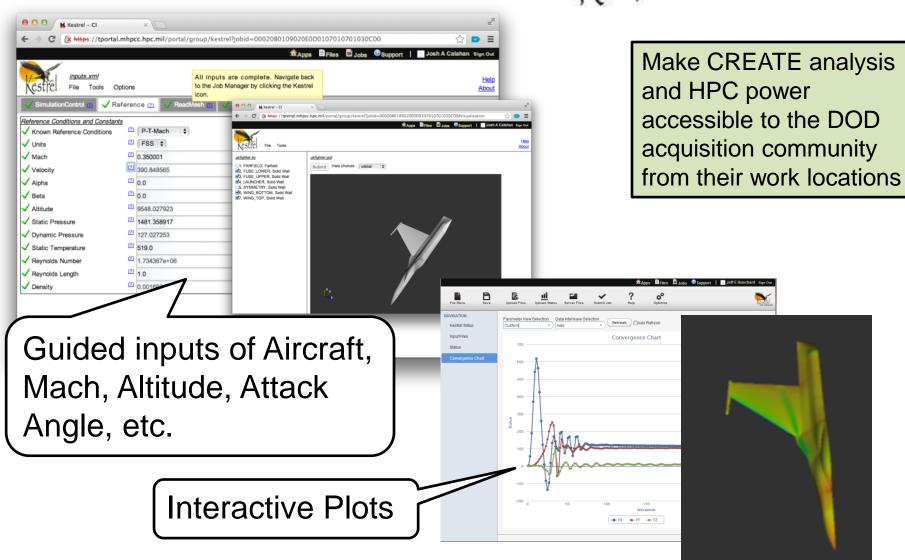


Apply the power of supercomputing to impact the mission!



### **HPCMP CREATE™ Kestrel**

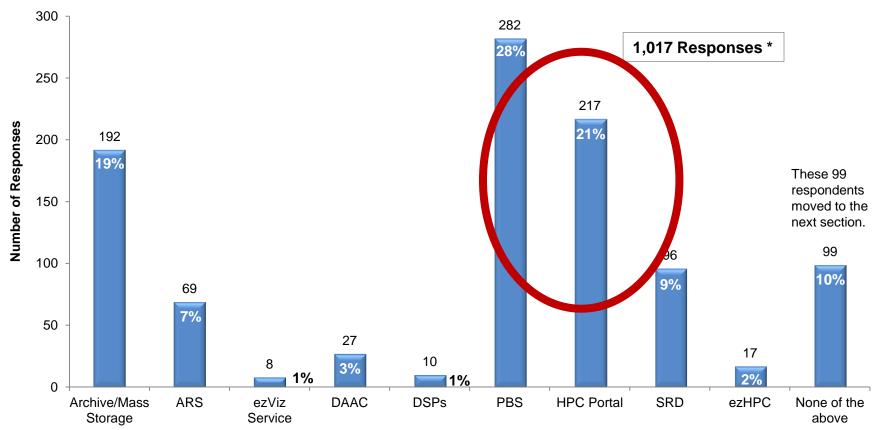
**Delivery Using the HPC Portal** 





## **User Satisfaction Survey FY14 Services, Tools and Capabilities**

Have you used/taken advantage of any of the following Program services, tools and/or capabilities? (choose all that apply\*)

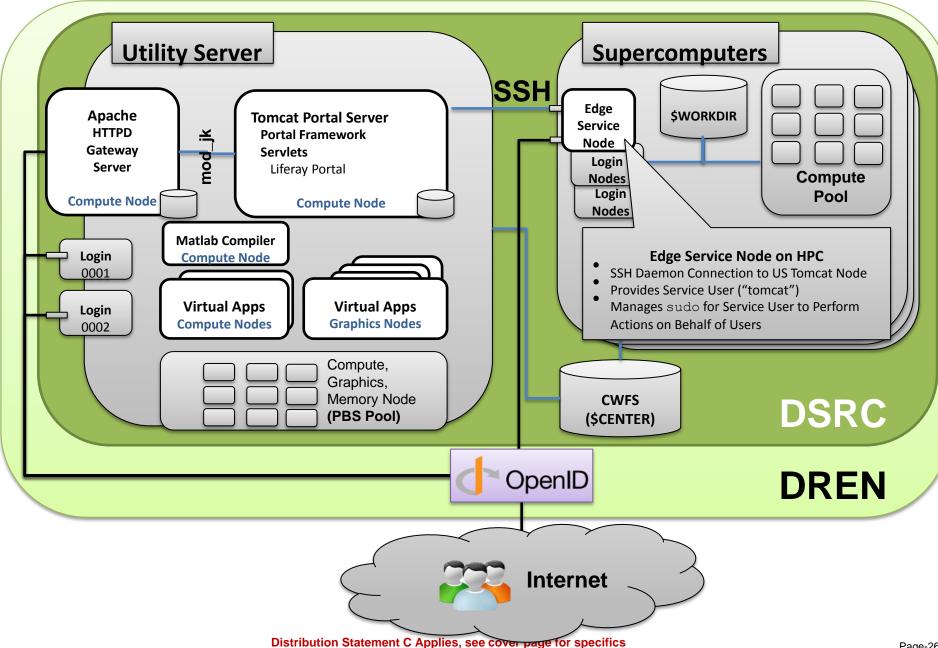


If only S/AAA role selected, respondent skipped this section.

<sup>\*</sup> Allowed for doubled counting.

#### **Portal Architecture**







#### **Advance Reservation Service**

(https://reservation.hpc.mil/)

- Advance Reservation Service (ARS) capability is available on all production systems
- ARS allows a user to reserve a block of time for their use on any system for which the user has an account
- This calendar gives a quick synopsis of how many nodes are available for an entire day on GARNET. For example, if 78 nodes are used for one how during the day, the number that will appear is 237. By clicking on each day a smallelle fluid, you can did flow to be seen blower and the process are using the first reached by the number that will appear is 237. By clicking on each day a smallelle fluid, you can did flow to be seen blower and energy dates are using the first reached and those and energy dates are using the first reached and the seen and the scholar remains available on the stabilities are using the first reached and the scholar remains a valuable on the stabilities.

  Monday Tuesday Wednesday Turnday Friday Saturday

  Monday Tuesday Wednesday Turnday Friday Saturday

  20 21 22 22 24 25 26

  21 22 22 24 25 26

  22 22 24 25 26

  23 23 23 23 22 22

  24 25 26

Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
			16 269	17 239	18 239	19 239
20	21	22	23	24	25	26
303	303	315	0	315	299	299
27 299	28 299	29 315	30 315	31 315	1 315 FEBRUARY	2 315
3	4	5	6	7	8	9
315	315	315	315	315	315	315
10	11	12	13	14	15	16
315	315	315	315	315	315	315
17	18	19	20	21	22	23
315	315	315	315	315	315	315

U.S. DEPARTMENT OF DEFENSE

Reservation Reques

e HPCMP is now accepting reservations for batch time on 315 nodes of GARNET (10080 cores @ 32 resper node). Renervations of 8 hours or less are limited to shif of this adjustation (157 nodes). A maximum 157 nodes(5040 cores) per reservation are available on all GARNET reservations. Advance reservations by be requested on GARNET for any available dates up to SEP <sub>20,0</sub> 2013 at 23:95 CDT. Upon reservation fring, you will have sole access to reserved nodes for the duration of your requested time periods.

- A reservation can be made up to 24-hours in advance of the requested time slot for up to 8 wall-clock hours of usage
- A reservation can be made up to 3 months in advance for a block of up to 168 wall-clock hours
- A single user can reserve up to 50% of ARS resources available on a system
  - Typically, 25% of a system is available for ARS



## **Dedicated Support Partition (DSP)**

- Intended to support large computational work that would benefit from a dedicated system partition and not be addressable with advance reservation service or utility server
  - Long-term recurring work (30–360 days)
  - Developmental testing

#### AFRL DSRC

- 57K Xeon Ivy Bridge Cray Cluster, 2 DSPs, partition sizes 512 and 5000 cores
- 73K Xeon Sandy Bridge SGI Cluster, 1 DSP, partition size 1088 cores

#### ARL DSRC

- 20K Xeon Sandy Bridge IBM Cluster, 1 DSP, partition size 1040 cores
- 17K Xeon Sandy Bridge IBM Cluster, 1 DSP, partition size 1088 cores

#### ERDC DSRC

150K AMD Interlagos Cray Cluster, 1 DSP, partition size 1184 cores

#### MHPCC DSRC

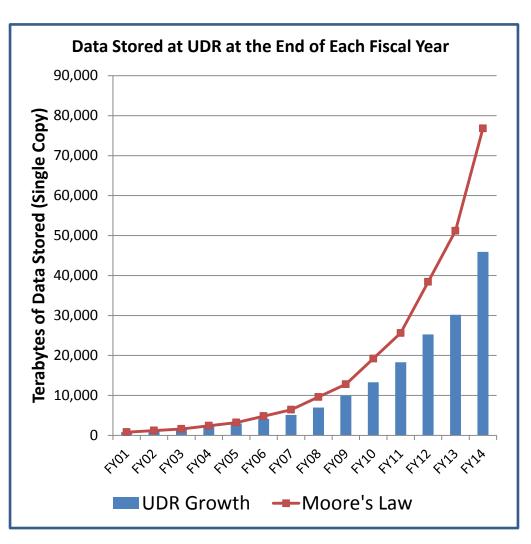
- 12K Xeon Sandy Bridge IBM Cluster
- 6 DSPs, with partition sizes ranging from 32 cores to 1024 cores

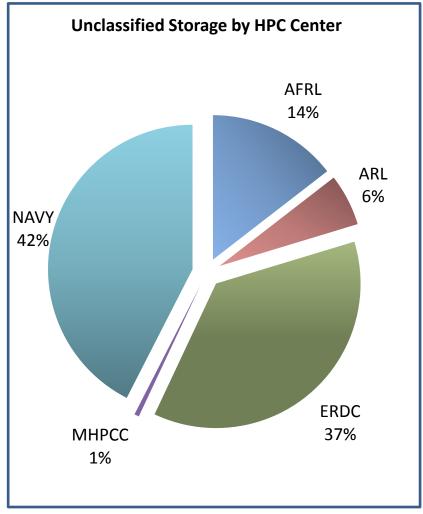
#### Navy DSRC

19K Xeon Sandy Bridge IBM Cluster, 1 DSP, partition size 1024 cores



## **Long-Term Data Storage**







### **Dedicated HPC Project Investments (DHPIs)**

Dedicated HPC Project Investments (DHPIs) are modest-sized HPC systems awarded to technically-sound, mission-critical projects that cannot be performed with shared HPCMP resources due to classification level, special operational requirements (e.g., real-time turnaround) and/or use of emerging technologies

#### **Awarded DHPI Proposals FY14** (\$4.1M)



Real-Time Hadoop/Hbase High Performance Computing System (\$0.35M) Army/ATEC, WSMR, NM and Aberdeen Test Center, MD



Applications of the Digital Thread to Capability Planning & Analysis (\$2.25M) Air Force/AF Lifecycle Management Center, Wright-Patterson AFB, OH



On-Demand GPU Applications (\$1.5M) Navy/Naval Research Laboratory, DC

9 proposals received: 3 selected, 1 recommended for DSRC support TI BOA available for system acquisition, or obligate by 30 Sept



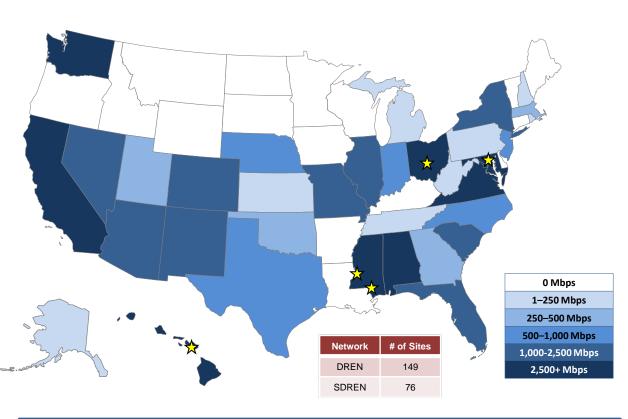
## Networking

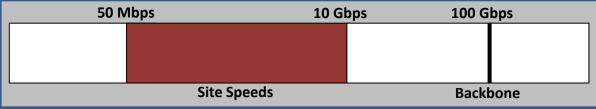
### **Defense Research & Engineering Network (DREN)**

- DoD's premier RDT&E network

   focused on S&T and T&E –
   provides separation from

   Warfighting networks
- High-bandwidth, low-latency, full-service network
- Connects DoD Supercomputing Resource Centers (DSRCs) and users
- Secret-level network overlay on DREN backbone (SDREN)
- DREN III provides 50 Mbps to 10 Gbps service to DoD sites; across a 100 Gbps backbone
- Fully supports IPv6 & Multicast
- Enables research and test events requiring low-jitter, lowlatency and high-throughput
- One component of the Global Information Grid – Sister network to NIPRNet/SIPRNet







## **DREN III Highlights**

- DREN III contract successfully awarded in late 2012
  - CenturyLink received a 10 year IDIQ award, worth up to \$750M
  - New technical capabilities such as Ethernet and Optical Services
  - Significantly increased capability

Attribute	DREN II	DREN III	Ratio
Backbone (in Gbps)	10	100	~10
Aggregate On-ramp BW (in Gbps)	30	150	~5
Internet BW (in Gbps)	4	20	~5

Transition of 149 sites completed by 17June 2014



## Resource Management Requirements & Allocations

#### Mission

 To ensure HPCMP computational resources are appropriately distributed across the DOD RDT&E communities based upon documented requirements and Service priorities

#### Goals

- Gather, analyze, store, and report requirements, allocation and usage information on DOD computational projects and HPC systems performance to support DOD/Services/Agencies allocation, program planning, acquisition, and operational processes
- Execute and report customer community special interest activities, such as requirements site visits and Challenge and Frontier Project allocation processes

#### Requirements

Requirements Gathering and Site Visits

#### pIE

Portal to the Information Environment

#### **Allocations**

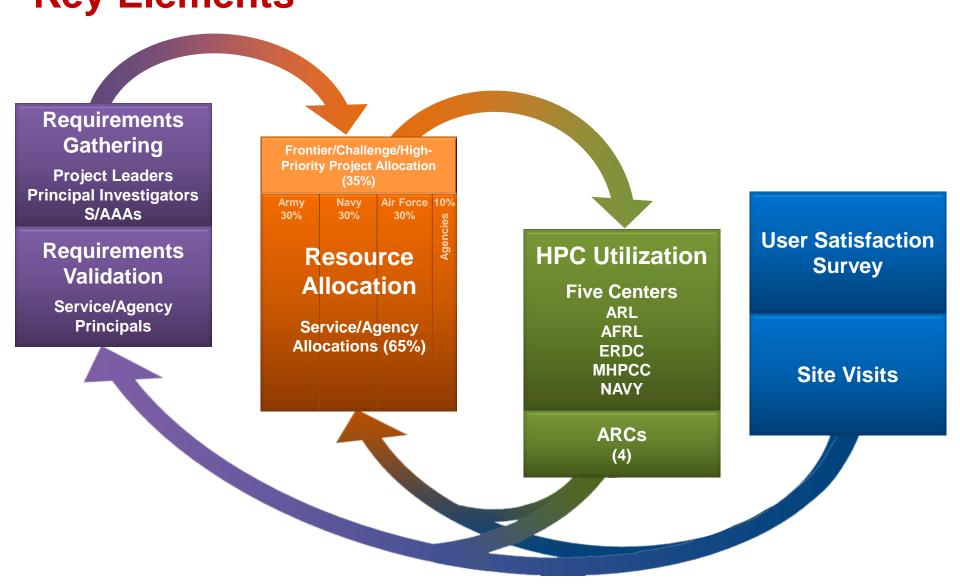
Challenge Projects, Frontier Projects, Capability Applications Projects (CAPs)

#### **Metrics**

**HPC Systems Performance** 

## **HPCMP** Resource Management Process **Key Elements**







## The Information Environment (IE) (and Its Portal (pIE))

- Authoritative database for all user, project, and subproject information
- All accounts must be approved by the Service organization's Service/Agency Approval Authority (S/AAA)
- Requirements questionnaire incorporated into IE
- All training requirements part of IE
  - Information Assurance tracked now
  - Additional training will be tracked with January release
- Constantly evolving to accommodate new Program initiatives
  - International Access forms
  - Unix Groups
  - New Software Request Forms
  - Dedicated Support Partitions
  - On-demand downloads available to Center Site Account Managers

## PPP WODERNIZATION PROGRAM

### **HPCMP Allocation Management Process**

- Allocate 100% of all resources at DoD Supercomputing Resource Centers
- Approximately 70% allocated to Services/Agencies for distribution
  - Resources are delegated to the Services/Agencies for allocation to Service/Agency mission priorities
    - 30/30/30/10 (Army/Navy/Air Force/Agencies) split
  - Mechanisms in place to support Service/Agency allocations for high priority, advance reservations, and dedicated workloads
- Frontier plus Challenge plus Service/Agency High-Priority is targeted at 35% of all resources allocated
  - Challenge Projects will end in FY 2015
  - FY 2016 target is for Service/Agency High-Priority and Frontier to share the 35% pool



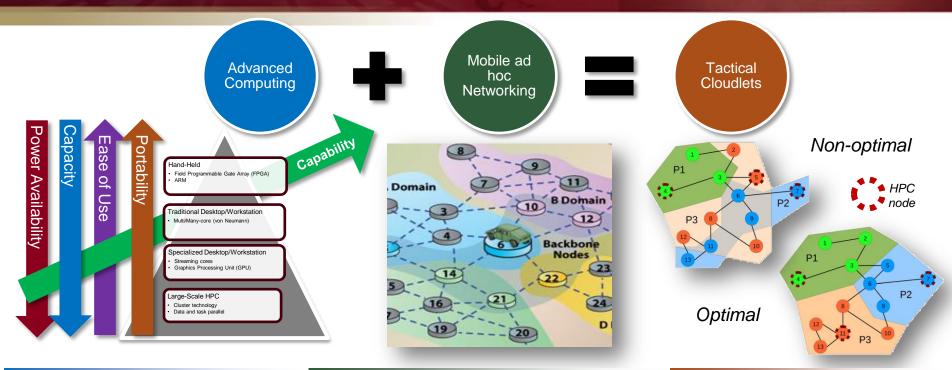
## **HPC Systems Performance Metrics**

- HPC Systems Performance Metrics are aggregated across all Program systems and sent to the HPC community monthly
- High-level metrics include:
  - System usage
  - Turnaround time (expansion factor)
  - Value metric combines system usage and turnaround time into one overall metric which characterizes the total value of the computational time to the user community
- New metrics being developed to track use and efficiency of advance reservation service



#### **Tactical Cloudlets**





- Focus on performance tailored per core
- No wasted processing cycles
- Power-aware and powerefficient
- Packing High
   Performance Computing
   (HPC) in small footprints

- Heterogeneous devices (and links)
- Hostile environments (short term, mobile, limited bandwidth)
- Cloud services too costly
- Virtually impossible to acquire real-time processing (via HPC)

- Efficiently distributed HPC nodes (seeding)
- Balanced processing capacity with demand
- Limiting hop distance between client-server